Unlocking Legal Efficiency Gains in School Switch Programs

Andersson, Tommy; Lager, Nils; Kessel, Dany

2026

Document Version:
Other version

Link to publication

Citation for published version (APA):
Andersson, T., Lager, N., & Kessel, D. (2026). *Unlocking Legal Efficiency Gains in School Switch Programs*. (Working Papers; No. 2026:4).

Total number of authors:
3

Working Paper 2026:4

Department of Economics
School of Economics and Management

Unlocking Legal Efficiency Gains in School Switch Programs

Tommy Andersson
Nils Lager
Dany Kessel

May 2026

1666

LUND
UNIVERSITY

Unlocking Legal Efficiency Gains in School Switch Programs

Tommy Andersson[†], Nils Lager[‡], and Dany Kessel[§]

May 21, 2026.

Abstract

We study school switches in which students who have already been assigned seats request reassignment. Unlike in regular school choice, stability need not be legally binding in such programs, since students' priority-based claims have already been satisfied upstream. Efficiency then naturally becomes the primary design objective. To quantify the potential legal efficiency gains, we compare Deferred Acceptance (DA) with Top Trading Cycles (TTC) and an efficiency-adjusted DA (EADAM) using Swedish data. TTC and EADAM triple switching rates compared to DA. This difference is driven by the absence of slack in effective capacity and shrinks rapidly when introducing additional seats.

JEL Classification: C78, D47, D78.

Key words: School switch, Efficiency, Stability, Matching algorithms, Capacity slack.

1 Introduction

Theoretical work on school choice has long emphasized the trade-off between stability and efficiency (Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu et al., 2009). Empirically, stable mechanisms, like the (student-proposing) Deferred Acceptance (DA) algorithm, are far more widely used than efficient mechanisms, like the Top Trading Cycles (TTC) mechanism (Abdulkadiroğlu and Andersson, 2023). A key reason is that laws and regulations are often interpreted as requiring priority-based stability, even when doing so yields efficiency losses. This perspective is formalized by Ehlers and Morrill (2019), who define a notion of (il)legal assignments in school choice. Priority constraints can be treated as externally imposed legal requirements that restrict the set of admissible matchings, even when violations would yield Pareto improvements. Under

---

[†] Department of Economics, Lund University. Email: `tommy.andersson@nek.lu.se`.

[‡] Department of Economics, Stockholm School of Economics. Email: `nils.lager@hhs.se`

[§] Svensk Utbildningsteknologi AB. Email: `dany.kessel@mittskolval.se`

this interpretation, stability is binding precisely when priority-based legal claims are at stake, but priority violations that no feasible reassignment could redress carry no such claim.

In Sweden, for example, the legal basis for school choice programs rests on one paragraph in the law (*Skollagen*, ch. 9, §15). It mandates that students should be placed at their most preferred school unless doing so would override "another student's legitimate claim to placement at a school near their home" or cause "significant organizational or economic difficulty" for the school district. Courts have interpreted this as effectively requiring priority-based stability, where proximity determines priority, along with sufficient (not necessarily unlimited) school capacities in the district.

This paper studies school switch programs. These programs apply to students who have already been assigned to schools through a school choice program and who have a wish, but not a legal right, to be reassigned.[1] While institutional details vary across Swedish school districts, many of them use DA in both their school choice and school switch programs. We argue that there is an important legal and conceptual distinction between these two types of programs.

First, the primary rationale for stability, namely the protection of priority-based legal claims, has already been satisfied at the initial assignment stage through the school choice program. Second, because participation in school switch programs is voluntary, any priority violations that arise do not overturn previously established rights, but instead reflect mutually beneficial reallocations among students who already hold seats. The stability constraints are therefore no longer binding in the same way in school switch programs as they are in school choice programs.[2] Once stability is no longer is legally required, efficiency naturally becomes the primary design objective.

To quantify the potential (and legal) efficiency gains by transferring from stable to efficient mechanisms in school switch programs, we use school switch data between 2020 and 2026 from six Swedish school districts. They all employ DA for both school choice and school switches, and our counterfactual analysis computes outcomes under TTC and the Efficiency-Adjusted Deferred Acceptance (EADAM) algorithm of Kesten (2010) as implemented by Faenza and Zhang (2022) with fully consenting students. The results indicate that only around 6 percent of participating students are part of a school switch under DA, while TTC and EADAM each roughly triple that rate. Because the DA-to-TTC and DA-to-EADAM efficiency gaps are nearly identical, the efficiency loss reflects DA's insistence on stability rather than any feature specific to either Pareto efficient algorithm.

The quantified efficiency gains in our data are substantially larger than those typically reported in school choice settings. For example, Abdulkadiroğlu et al. (2009) and Abdulkadiroğlu et al. (2017) report modest increases in placements at higher-ranked schools under TTC com-

---

[1] This excludes year-round processes that address families' legal right to urgent placements, for example, due to moving into the district, moving within the district, or relocation due to bullying.

[2] Even school districts that aim to preserve stability face challenges over the school year, since switches are inherently dynamic, with applications and seat availability evolving over time, see Doval (2022).

pared to DA. We demonstrate that this difference can potentially be explained by the absence of slack in effective capacity in the considered setting. In school switch programs, each participating student is effectively endowed with one school seat, so any reassignment necessarily displaces another student, making stability constraints bind tightly and sharply limiting DA's ability to reallocate. By contrast, in standard school choice environments, excess capacity and outside options imply that many assignments can be adjusted without violating priorities, leaving less room for improvement. In this respect, school switches are closer to settings studied in teacher assignment than to school choice.[3] In teacher assignment, the number of positions is close to the number of applying teachers, and relaxing stability yields efficiency gains of the same magnitude as in this paper (see Combe et al., 2022).

2 Model and Data

We start by providing a simple description of a school switch program. For a given school district and year, all schools in the district and all students who participate in the school switch program are collected in the sets S and I, respectively. Each student $i \in I$ has been assigned to a school $s_i \in S$ in the regular school choice program. In a given year, student i enters the school switch program at date $t_i \in \{1, \ldots, 365\} \equiv T$ and submits a rank-ordered list (ROL) $\mathcal{L}_i$ over the schools in S that are strictly preferred to s_i. School priorities over students are determined by a single tie-breaking lottery[4] (since priorities have already been respected upstream in the school choice program) subject to that student i have the highest priority to a seat at school s if $s = s_i$.[5]

Every k days, the school district runs the school switch program among all active students, that is, students who have entered the program by that date but have not been part of a switch in any earlier round. The capacity of each school $s \in S$ at a given run is given by the number of active students currently assigned to s, i.e., the number of students with $s_i = s$ who remain active in the program. Students who are matched exit the program, while unmatched students carry over to the next run.

We use school switch data from 2020 to 2026 for six Swedish school districts (municipalities). For each district, year, and student i, we observe S, I, t_i, and $\mathcal{L}_i$. The frequency with which the school switch program is run varies across districts, ranging from 2 to 6 times per year (see Table 1). To isolate the effect of the mechanism used for school switches (DA, TTC, and EADAM), we rerun the program using the available data for run frequencies $k \in \{10, 30, 90, 180, 365\}$. A complication is that application windows differ across districts, so not all students are allowed to apply at arbitrary dates in T (see Table 1). To facilitate comparability

---

[3] Note that school switch programs are structurally analogous to housing markets with existing tenants (Shapley and Scarf, 1974), since students already "own" seats at schools.

[4] Our results are robust to running the simulations with each municipality's original priorities.

[5] The student's current school is added to the bottom the student's ranked-ordered list, so students always can "bounce back" to their current school in case of rejection.

across specifications, we randomize application dates within the relevant intervals to simulate uniform arrivals over the year.[6]

School district (municipality)	Yearly runs	Application window
Linköping	3	14 days
Malmö	2	21 days
Nacka	6	7-14 days
Norrköping	4	16 days
Uppsala	6	Year-round
Växjö	2	Varying from year to year

Table 1: Run frequencies and application windows for the school districts.

3 Empirical Results

Table 2 summarizes the results for run frequency $k = 90$. Aggregated across the six school districts, only about 5–6 percent of participating students are part of a school switch under DA, whereas TTC and EADAM increase this rate to about 15–23 percent, depending on school district. This is mirrored by a near-identical increase in first-choice placements. Interestingly, EADAM and TTC track each other very closely (always within one percentage point) and neither mechanism strictly dominates the other. The overall pattern is remarkably consistent across districts.

As shown in Figure 1, the same pattern holds across all considered run frequencies. DA improves modestly with more frequent runs, as additional rounds create opportunities for later-arriving students to switch before being blocked by stability constraints. However, these gains remain small relative to those achieved by the efficient mechanisms at any given k, and the gap between DA and TTC/EADAM persists throughout.

We next demonstrate that the efficiency gap between DA and TTC/EADAM can potentially be explained by the absence of slack in effective capacity in school switch programs. In other words, the gains from replacing DA are themselves sensitive to another institutional feature: How tightly capacities bind. To make this point, suppose that we add x percent additional (empty) school seats per active student, where x is defined as the total number of additional seats divided by the number of active students. In our baseline model with no slack, $x = 0$. As shown in Figure 2, the gap between DA and TTC/EADAM is large when capacities are tight (i.e., when x is low),

---

[6] More formally, suppose that the program is run at dates k and $2k$, and consider the dates $\{k, \dots, 2k-v, \dots, 2k\}$. Assume that the application window for students entering the program at date $2k$ is $\{2k - v, \dots, 2k\} \equiv \hat{T}$, so that $t_i \in \hat{T}$ in the data for any student who is part of their first switch round at date $2k$. In our analysis, we instead draw t_i uniformly from $\{k + 1, \dots, 2k\}$, which allows us to also study switch rounds at dates $\{k + 1, \dots, 2k - v - 1\}$.

School district	**No. ROLs**	**Percent switchers**			**Percent 1st Choices**		
		DA	EADAM	TTC	DA	EADAM	TTC
Linköping	3,995	6.23	20.92	20.39	4.25	15.66	15.73
		(0.57)	(0.23)	(0.22)	(0.40)	(0.24)	(0.23)
Malmö	6,922	5.19	22.65	21.83	3.12	15.49	15.79
		(0.49)	(0.27)	(0.26)	(0.33)	(0.25)	(0.23)
Nacka	3,683	5.60	20.43	19.89	4.34	17.37	17.43
		(0.60)	(0.24)	(0.22)	(0.51)	(0.24)	(0.22)
Norrköping	2,878	7.14	23.35	22.95	5.71	19.55	19.70
		(0.76)	(0.38)	(0.37)	(0.64)	(0.34)	(0.34)
Uppsala	4,286	4.57	19.30	18.68	3.07	14.59	14.71
		(0.62)	(0.27)	(0.26)	(0.44)	(0.24)	(0.24)
Växjö	592	6.41	15.33	15.45	5.36	13.64	13.75
		(1.25)	(0.36)	(0.32)	(1.16)	(0.38)	(0.35)
Aggregated	*22,356*	*5.61*	*21.23*	*20.62*	*3.91*	*16.13*	*16.29*
		(1.01)	*(1.74)*	*(1.66)*	*(1.04)*	*(1.62)*	*(1.60)*

Table 2: Summary statistics by school district (run frequency $k = 90$). Means and standard deviations across 250 Monte Carlo draws (standard deviations in parentheses).

and shrinks rapidly as schools are given additional slack. With sufficiently many additional seats, roughly 80 percent of all students are part of a switch for all three mechanisms.

As Kesten (2010) shows, DA's inefficiency is determined by the number of interrupting pairs in the matching.[7] Although a formal characterization is left for future work, the intuition behind the shrinking efficiency gap is clear: An interrupting pair can only arise at an over-subscribed school. Adding capacity slack mechanically reduces the set of over-subscribed schools, and with it the potential for interrupting pairs.

Acknowledgements

Tommy Andersson gratefully acknowledges financial support from the Jan Wallander and Tom Hedelieus Foundation (grant no. P22–0087).

---

[7] An interrupting pair is a student-school pair in which a student temporarily displaces a lower-priority student from a school during DA's execution, but is herself subsequently rejected from that school, so the displacement is wasted.

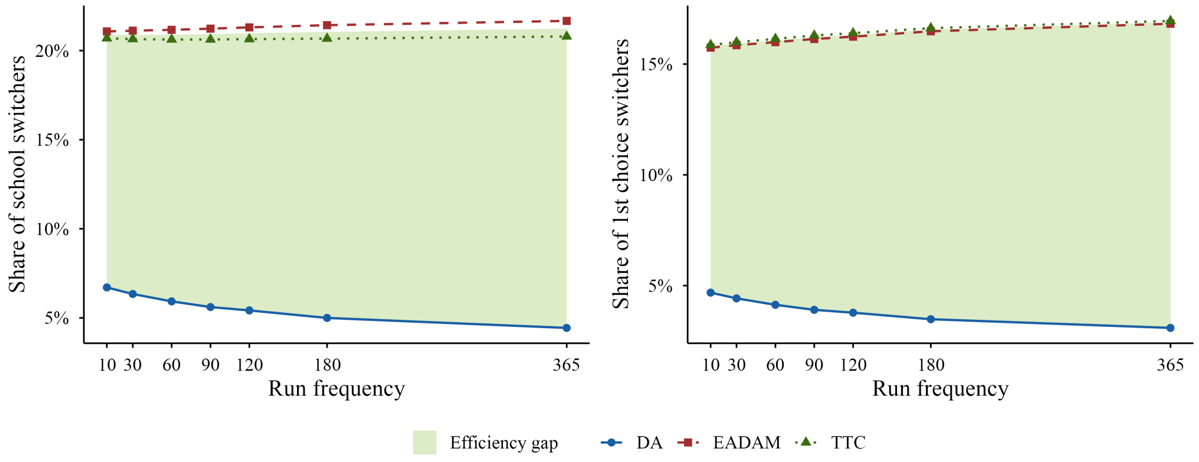

Figure 1: Outcomes by run frequency (days). Means across 250 Monte Carlo draws, applicant-weighted across the six school districts. The shaded band represents the gap between DA and the TTC/EADAM mean.

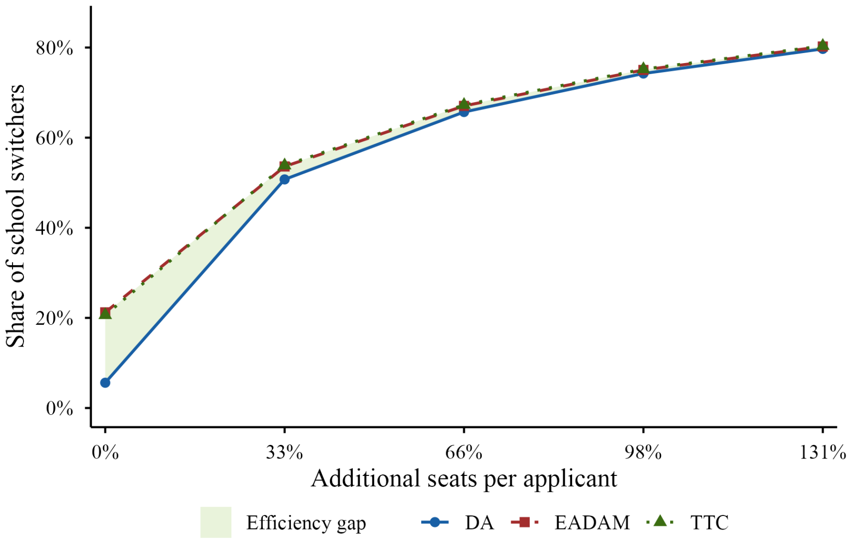

Figure 2: Share of students involved in a switch for run frequency $k = 90$ with base capacity added uniformly across all schools. Means across 250 Monte Carlo draws, applicant-weighted across the six school districts. The shaded band represents the gap between DA and the TTC/EADAM mean.

References

Atila Abdulkadiroğlu and Tommy Andersson. School Choice. In E. A. Hanushek, S. Machin, and L. Woessmann, editors, *Handbook of the Economics of Education*, volume VI, chapter 3, pages 135–185. Elsevier, Amsterdam, 2023.

Atila Abdulkadiroğlu and Tayfun Sönmez. School Choice: A Mechanism Design Approach. *American Economic Review*, 93:729–747, 2003.

Atila Abdulkadiroğlu, Parag A. Pathak, and Alvin E. Roth. Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match. *American Economic Review*, 99(5):1954–1978, 2009.

Atila Abdulkadiroğlu, Nikhil Agarwal, and Parag A. Pathak. The Welfare Effects of Coordinated Assignment: Evidence from the New York City High School Match. *American Economic Review*, 107(12):3635–3689, 2017.

Julien Combe, Olivier Tercieux, and Camille Terrier. The Design of Teacher Assignment: Theory and Evidence. *The Review of Economic Studies*, 89(6):3154–3222, 2022.

Laura Doval. Dynamically Stable Matching. *Theoretical Economics*, 17(2):687–724, 2022.

Lars Ehlers and Thayer Morrill. (Il)legal Assignments in School Choice. *The Review of Economic Studies*, 87(4):1837–1875, 2019.

Yuri Faenza and Xuan Zhang. Legal Assignments and Fast EADAM with Consent via Classic Theory of Stable Matchings. *Operations Research*, 70(3):1873–1890, 2022.

Onur Kesten. School Choice with Consent. *The Quarterly Journal of Economics*, 125(3):1297–1348, 2010.

Lloyd Shapley and Herbert Scarf. On Cores and Indivisibility. *Journal of Mathematical Economics*, 1(1):23–37, 1974.